Remarks

Reconsideration and withdrawal of the final Office Action, and allowance of all pending claims are respectfully requested. Claims 50-53, 56, 57, 59, 61-63, 66-69, 72-74, 76-81, 84-86 & 88-90 remain pending.

All pending claims (i.e., claims 50-53, 56, 57, 59, 61-63, 66-69, 72-74, 76-81, 84-86 & 88-90) stand rejected under 35 U.S.C. §102(e) as being anticipated by Snaman, Jr., et al. (U.S. Patent No. 6,243,744 B1; hereinafter Snaman, Jr.). This rejection is respectfully traversed, and reconsideration and withdrawal thereof are requested.

Shared Nothing:

The present invention relates to a processing protocol for managing processing groups of a *shared nothing* distributed computing environment. A shared nothing distributed computing environment is a particular computing environment architecture that is well known to one of ordinary skill in the art. By way of example, an internet search for a *shared nothing system* uncovers many references to the particular architecture at issue. Appendix A and Appendix B are two examples of parallel hardware architecture descriptions of a shared nothing system. As noted in these materials, a shared nothing system is a particular system architecture wherein only one processing node is connected to a given disk, with the *nodes being interconnected via a network only*. This particular architecture is referred to as *shared nothing* since there is no shared disc between two more processing nodes of the system. A *shared nothing* computing environment is a distinct computing architecture from a *shared disk* architecture.

Shared Disk:

A *shared disk* system refers to a computing architecture wherein one or more disks are shared by two or more CPUs or nodes in the system. In a shared disk system, the two or more nodes sharing the one or more disks have access to the same disks. Central to the shared disk system is the concept that one or more disks are shared by two or more nodes of the system.

In the instant case, Applicants' recited invention comprises protocol for a shared nothing distributed computing environment, while the applied patent (Snaman, Jr.) presents protocol for a shared disk system. This difference is significant, and based thereon, Applicants respectfully submit that there is no anticipation of the recited invention based upon the teachings of Snaman, Jr.

Snaman, Jr. describes a technique for sharing a resource among a cluster of devices in a computer network. As shown in FIGS. 1A-1E & 4A-4G, Snaman, Jr. discloses a computing environment wherein a *quorum disk* is employed in forming a cluster by voting devices attempting to reach a quorum. The *quorum disk* described by Snaman, Jr. is a non-processing voting device that is *shared by at least two nodes of the environment*. Because the quorum disk is shared, the computing environment of Snaman, Jr. is a *shared disk* computing environment, not a shared nothing computing environment, as recited in Applicants' independent claims.

In this regard, Applicants respectfully traverse the characterizations contained at pages 12 & 13 of the final Office Action regarding the meaning of a *shared nothing* computing environment. These conclusions are separately addressed below.

At page 12, lines 19 & 20 of the final Office Action, it is stated that:

> In response to A. above, it is the teaching of the "quorum disk" that in fact teaches that the system of Snaman, Jr. is a shared nothing computing environment.

As noted above, the quorum disk in Snaman, Jr. is a *shared disk* that is taught to be shared by nodes A & B of the system. This shared disk is shown in FIGS. 1A-1E, as well as FIGS. 4A-4G. In each of these figures, two nodes connect to the quorum disk by dedicated lines 20, 22 (in FIGS. 1A-1E) or dedicated lines 100, 102 (in FIGS. 4A-4G). Because nodes A & B share the quorum disk in Snaman, Jr., Snaman, Jr. presents a *shared disk* system. As noted above, a *shared nothing* computing environment means that there is not a shared device or disk between the nodes of the environment. This interpretation of Snaman, Jr. and Applicants' independent claims is believed well supported by the art, as represented by Appendices A & B attached herewith. The difference in architecture is significant since the protocols employed in the architectures are vastly different.

At page 12, line 20 – page 13, line 3 of the final Office Action, it is stated:

> It is well known in the art that a shared nothing computing environment is a computing environment wherein applications cannot be distributed across multiple nodes.

This conclusion is believed a mischaracterization of the term "shared nothing computing environment". The term "shared nothing computing environment" refers to a particular hardware architecture, as explained above. Further, it is incorrect to state that applications cannot be distributed across multiple nodes in a shared nothing computing environment since they can. In a shared nothing computing environment, the only connection between nodes is the network itself. However, there are existing protocols (such as MPI or DB2) which allow applications to be distributed across multiple nodes in a shared nothing distributed computing environment.

At page 13, lines 2-5 of the final Office Action, it is stated:

> As such, the reason that it is possible for a node to take over running an application when the active node fails in a shared nothing computing environment, is because all the nodes in the cluster are connected to a shared storage mechanism (quorum disk).

As noted, a shared nothing computing environment refers to a particular computing environment architecture wherein there is no shared device or disk between two or more nodes of the environment. Applicants recite a protocol for managing processing groups within such a shared nothing computing environment (i.e., an environment where there is no shared quorum disk).

At page 13, lines 5-6 of the final Office Action, it is stated that:

> In fact, quorums are only used in a shared nothing environment.

Applicants respectfully traverse this statement. As taught by Snaman, Jr., quorum is employed in a shared disk system as well. The difference is in the protocol which employs the quorum. In Snaman, Jr., the quorum disk controls which nodes are granted quorum, while in a shared nothing computing environment, quorum is decided among the nodes using a protocol such as recited by Applicants.

At page 13, lines 7-10 of the final Office Action, it is stated:

> Quorum is necessary in a shared nothing clustering environment so that it is able to tell which node is active and which node or nodes are in stand by (to maintain consistent functionality), and also so that when failure occurs, only the partition that owns the quorum remains running the application (to maintain data consistency).

Applicants agree with this particular statement; however, note that in a shared nothing computing environment, the nodes can only count on communication via the network, whereas in a shared disk environment such as taught by Snaman, Jr., at least some of the nodes are also in direct communication via the shared disk. In the Snaman, Jr. protocol, the shared disk is employed as part of the protocol for responding to a failure.

At page 13, lines 10-12 of the final Office Action, it is stated that:

> The fact that all the nodes share a quorum disk does not teach away from a shared nothing computing environment but rather teaches of a shared nothing computing environment.

This statement is clearly contradictory to the well established meaning of *shared nothing* computing environment versus a *shared disk* computing environment. A shared disk computing environment means that there is a shared device (i.e., a disk) between two or more nodes of the computing environment. This shared disk plays a significant role in overseeing the establishing of a quorum subsequent to failure of a node in a shared disk computing environment. In contrast, the *shared nothing* computing environment has no shared device or shared disk between two more nodes of the computing environment. The only connection between nodes in a shared nothing computing environment is the network itself.

In addition to the above, Applicants respectfully submit that the protocol of Snaman, Jr. would be inoperable in a shared nothing computing environment since the shared disk protocol described therein relies upon the existence of the disk itself. A shared quorum disk protocol such as described by Snaman, Jr. pre-existed Applicants' invention and is directed to a different computing architecture than Applicants' recited invention. In a shared nothing computing environment, there is typically no physical way to connect a shared device, for example, between remote clusters of nodes. A shared disk environment such as described by Snaman, Jr. requires dedicated connections. These special hardware connections couple two or more nodes of the system to the shared disk, as illustrated in all of the Snaman, Jr. architecture figures. The

Snaman, Jr. shared disk computing environment is referred to in the art as a "twin-tailed" disk configuration, wherein two nodes of the system connect to the shared disk. This is not the environment recited by Applicants in the independent claims presented. In Applicants' independent claims, a shared nothing computing environment is expressly recited.

Given this characterization, Applicants respectfully submit that there is no anticipation of their recited invention based upon the teachings of Snaman, Jr., nor would Applicants' invention have been obvious to one of ordinary skill in the art based upon the teachings of Snaman, Jr. The protocol described by Snaman, Jr. expressly requires the existence of the shared disk. Further, one skilled in the art would not have modified the teachings of Snaman, Jr. so as to remove the shared quorum disk. Snaman, Jr. relies on the existence of the shared device for the clustering protocol described therein. The quorum disk in Snaman, Jr. is critical to situations in which no cluster has a majority, and the quorum disk provides the configuration information. Thus, once a processing node arrives at the quorum disk, the configuration can be obtained from the quorum disk. Without the quorum disk, the Snaman, Jr. protocol would be unworkable.

In contrast, Applicants' recited protocol does not rely upon nor implement shared disk protocol. In Applicants' recited invention, the environment is a *shared nothing* computing environment, meaning that there is not a shared device between the nodes of the environment (as represented by the Appendices A & B materials, as well as the understanding of one of ordinary skill in the art).

For at least the above-noted reasons, Applicants respectfully submit that the pending independent claims patentably distinguish over the teachings of Snaman, Jr. As such, reconsideration and withdrawal of the final rejection is respectfully requested.

The dependent claims are believed allowable for the same reasons as the independent claims, as well as for their own additional characterizations. For example, dependent claims 52, 68 & 80 recite quiescing activity that may effect the state prior to the updating. The term "quiescing" is a term of art which refers to the gradual quieting of a system or network. Quiescing occurs when a signal is sent to stop processing, and then a period of time passes while processes (gracefully) discontinue processing. The processes are essentially allowed to save their current state and exit voluntarily. In Snaman, Jr., column 13, lines 36-39, it is noted that

Snaman, Jr. presents a method and apparatus for stopping a partitioned cluster such that processing operations are not performed which would destroy cluster information on a shared resource. As explained at column 12, line 62+, Snaman, Jr. stops by crashing the cluster to remove any possibility of performing data processing operations that would destroy cluster information on the shared resource. This crashing of the cluster is clearly not an analogous to Applicants' recited quiescing activity that may effect state prior to the updating. The two are distinct functions, and one does not suggest the other. Because Snaman, Jr. is a quorum disk based approach, Snaman, Jr. cannot risk data corruption, and the cluster must crash. In contrast, Applicants' invention is not quorum disk based (i.e., is a shared nothing based protocol), and the processes are allowed to quiesce voluntarily.

In view of the above, Applicants respectfully request issuance of an indication of allowance for all pending claims.

Should the Examiner wish to discuss this application with Applicants' undersigned attorney, the Examiner is invited to contact their representative at the below-listed telephone number.

Respectfully submitted,

Kevin P. Radigan, Esq.
Attorney for Applicants
Registration No.: 31,789

Dated: December 01, 2006.

HESLIN ROTHENBERG FARLEY & MESITI P.C.
5 Columbia Circle
Albany, New York 12203-5160
Telephone: (518) 452-5600
Facsimile: (518) 452-5579

Attachments: Appendix A (9 pages)
Appendix B (8 pages)